



US006357041B1

(12) **United States Patent**
Pingali et al.

(10) **Patent No.:** US 6,357,041 B1
(45) **Date of Patent:** Mar. 12, 2002

(54) **DATA-CENTRIC MULTI-LEVEL BLOCKING**

(75) Inventors: **Keshav K. Pingali**, Ithaca, NY (US);
Induprakas Kodukula, Dallas, TX
(US); **Nawaaz Ahmed**, Ithaca, NY (US)

(73) Assignee: **Cornell Research Foundation, Inc.**,
Ithaca, NY (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/450,888**

(22) Filed: **Nov. 29, 1999**

Related U.S. Application Data

(63) Continuation of application No. PCT/US98/10938, filed on
May 29, 1998.

(60) Provisional application No. 60/047,382, filed on Jun. 2,
1997.

(51) Int. Cl.⁷ **G06F 9/45**

(52) U.S. Cl. **717/9**

(58) Field of Search **717/9**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,175,837 A	*	12/1992	Arnold et al.	711/152
5,717,893 A	*	2/1998	Mattson	711/129
5,754,888 A		5/1998	Yang et al.	395/872
5,790,828 A		8/1998	Jost	395/404

OTHER PUBLICATIONS

Steve Carr and R.B. Lehoucq, "Compiler Blockability of
Dense Matrix Factorizations", Oct. 2, 1996.

Steve Carr and R.B. Lehoucq, Chapter I—A *Compiler-
Blockable Algorithm for QR Decomposition*. Oct. 1996.

Michael E. Wolf and Monica S. Lam, "A Data Locality
Optimizing Algorithm", Computer Systems Laboratory,
Stanford University, CA 94305. Jun. 1991.

Michal Cierniak and Wei Li, "Unifying Data and Control
Transformations for Distributed Shared-Memory
Machines", Department of Computer Science, University of
Rochester, Rochester, NY 14627. Jun. 1995.

Kathryn S. McKinley (University of Massachusetts at
Amherst), Steve Carr (Michigan Technological University),
and Chau-Wen Tseng (University of Maryland at College
Park), "Improving Data Locality with Loop Transforma-
tions". Jul. 1996.

Monica S. Lam, Edward E. Rothberg, and Michael E. Wolf,
"The Cache Performance and Optimizations of Blocked
Algorithms", Computer Systems Laboratory, Stanford Uni-
versity, CA 94305; Fourth Intern. Conf. On Architectural
Support for Programming Languages and Operating Sys-
tems (ASPLOS IV), Palo Alto, California, Apr. 9-11, 1991.
J. Ramanujam (Dept. of Electrical and Computer Engineer-
ing, Louisiana State University, Baton Rouge, LA
70803-5901) and P. Sadayappan (Dept. of Computer and
Information Science, The Ohio State University, Columbus,
OH 43210-1277), "Tiling Multidimensional Iteration
Spaces for Multicomputers". Oct. 1992.

(List continued on next page.)

Primary Examiner—Tuan Q. Dam

Assistant Examiner—John Q. Chavis

(74) *Attorney, Agent, or Firm*—Perkind, Smith & Cohen,
LLP; Christine M. Kuta; Jacob N. Erlich

(57)

ABSTRACT

A framework for improving program performance by
locality-enhancing transformations is presented. This frame-
work is appropriate for modern high-performance machines
that have a memory hierarchy. The invention orchestrates
the flow of data through the memory hierarchy directly, and
is thus able to overcome limitations of existing approaches.
This new approach allows for efficient execution of imper-
fectly nested loop programs which are ubiquitous in numeri-
cal calculations and database operations, and it can be
integrated into high-performance optimizing compilers.

12 Claims, 7 Drawing Sheets

